IBM POWER4 Chip Integration

John Petrovick and Ed Seymour
IBM, Enterprise Systems Group
Austin, Texas

Agenda

- POWER4 Overview
- Challenges
- Concurrent Design
- Floorplanning
- Timing
- Checking
- Summary

POWER4 Architecture Overview

- System On A Chip
- 2 Processors/Chip
- Integrated/Shared L2 Cache
- L3 Directory on chip, L3 Cache off chip
- High bandwidth server building block

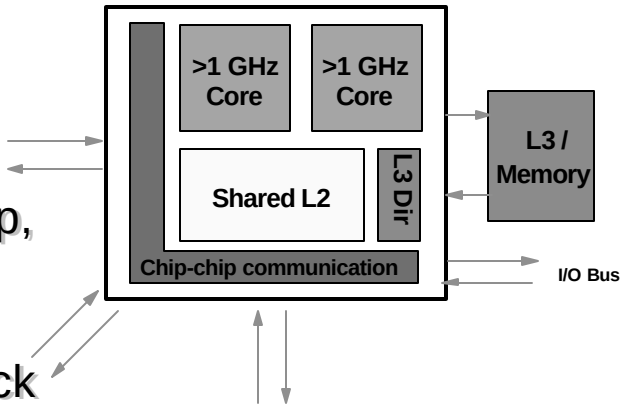

POWER4 Floorplan

- >1 GHz frequency
- 174M transistors
- 0.18um lithography
  - Silicon on Insulator (SOI)
- 7 levels of metal
  - Copper interconnect
  - >1 mile of wire
- 2208 signal I/O, >5500 total I/O
  - >1 Tb/sec pin bandwidth

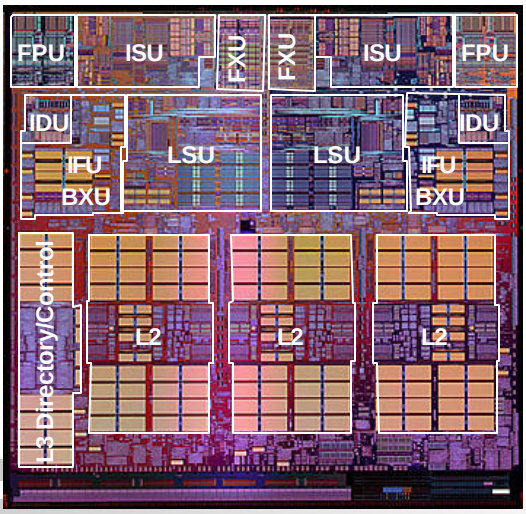

Challenges

- Quantum leap in integration complexity
  - 20M transistors to 170M+ transistors
- Dramatic increase in design complexity
  - Function, frequency
- Accommodate floorplan constraints
- Increased function + Higher frequency further complicates design
  - 3 clock cycles to cross the chip!!!
- Validate design
  - Function and performance

All while managing design turn-around-time!!!

Concurrent Design Approach

- Logic, timing, physical design optimized simultaneously within a unit
  - Unit/core/chip optimized in parallel
- Change management
  - EC isolation
  - Spare logic at all levels of hierarchy
- Flexible organization
  - n-Level physical and logical hierarchy
  - Physical != logical hierarchy

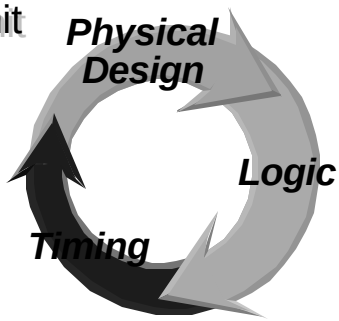

Concurrent Design Approach . . .

- Manage the complexity
  - Top-down budgeting with bottom-up feedback
- Enforce strict use of contracts
  - Timing
    - Arrival, departure times and clock phases
  - Physical
    - Size, pin location and layer, routing resources, fixed parent blockage
- High Level Design phase set area and timing contracts

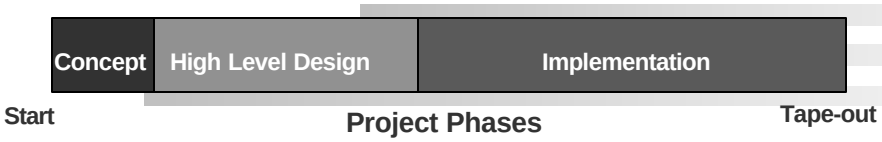

Floorplanning

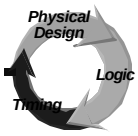

- Performance sensitive/noise critical and topologically difficult nets
- Pre-planned wiring as part of high level design
  - Concurrent with microarchitecture implementation
  - Exact track location, buffer and latch intervals
  - Wire width, space, hostility expectations
- Tied to timing methodology
  - Wirecodes for estimated timing based on different wiring topologies

Floorplanned Buses

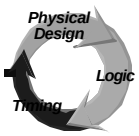

- Routing pre-specified
- Tools in place to automate process
  - ► Accommodates pre-routed buses
- Incorporated into common IBM Server Group design methodology

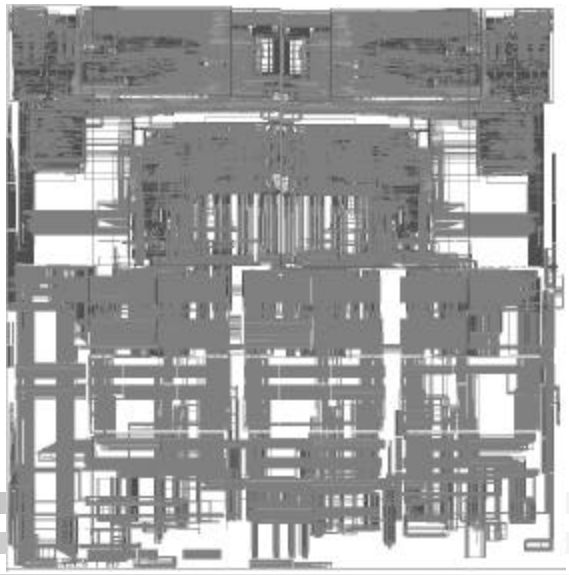

Timing

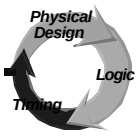

- Timing budgets throughout hierarchy set during High Level Design phase
- Topological wirecodes used to estimate final routed RC
- Minimized use of dynamic circuitry
- 3D Extraction at all hierarchy levels
  - ► Optimization at Custom/Unit/Core/Chip using boundary assertions
  - ► As design matured, flat core, then flat chip timing used for final refinement

Checking Strategy

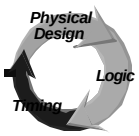

- Concurrent and hierarchical
  - Customs, units, control macros pre-verified
  - Build quality confidence up the hierarchy
- Use established physical contracts for "in-situ" checking
  - Cover used during routing
  - Fixed physical obstructions passed down the hierarchy

Trust *but* Verify

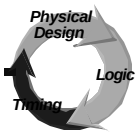

- Extensive methodology checks
  - Shapes inside boundary, correctly formed boundaries, on-grid pins, accessible pins, parent shapes touch only child pins, . . .
- Formal verification
  - Gold VHDL model vs. physical netlist
  - Hierarchy differences, buffer insertions validated

Summary

- Robust, flexible methodology and tools in place
  - Supports hierarchical/concurrent design and optimization for large complex system
  - Chip size stable 15 months prior to tape-out
- Supports rapid and frequent design turns
  - 12 Hours for full Netlist change, routing, and extraction
  - 16 Hours from VHDL change to tapeout
- Well positioned for higher frequency follow-on designs